

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: 28 October 2026

K. Talaulikar
A. MahendraBabu, Ed.
C. Filsfils
K. Swamy
Cisco Systems
S. Zandi
G. Dawra
LinkedIn
M. Durrani
Equinix
26 April 2026

BGP Link-State Extensions for BGP-only Networks
draft-ietf-idr-bgp-ls-bgp-only-fabric-05

Abstract

BGP is used as the only routing protocol in some networks today. In such networks, it is useful to get a detailed topology view similar to one available when using link state routing protocols. This document defines extensions to the BGP Link-state (BGP-LS) address-family and the procedures for advertisement of topology information in a BGP-only network.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 October 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. BGP Routing in the Fabric	3
3. Topology Collection Mechanism	4
3.1. Peering Models	5
4. Advertising BGP-only Network Topology	5
4.1. Node Advertisements	5
4.2. Link Advertisements	7
4.3. Prefix Advertisements	10
4.4. SR Policy Advertisements	11
4.5. SRv6 SID Advertisements	12
5. Procedures	13
5.1. Advertisement of Router's Node Attributes	13
5.2. Advertisement of Router's Local Links Attributes	14
5.3. Advertisement of Router's Prefix Attributes	16
5.4. Advertisement of Router's SR Policy Candidate Path Attributes	17
5.5. Advertisement of Router's SRv6 SID Attributes	17
6. Usage of BGP Topology	18
6.1. Topology View for Monitoring	18
6.2. SR-TE in BGP Networks	18
7. IANA Considerations	20
8. Security Considerations	21
9. Acknowledgements	21
10. References	21
10.1. Normative References	21
10.2. Informative References	23
Authors' Addresses	25

1. Introduction

Network operators are going for a BGP-only routing protocol for certain networks like Massively Scaled Data Centers (MSDCs). [RFC7938] describes the requirement, design and operational aspects for use of BGP as the only routing protocol in MSDCs. The underlying link and topology information between BGP routers is abstracted in this design for improving scalability and stability in a large scale

network. As a result, a detailed topology view consisting of nodes, links and prefixes that is available when operating link-state routing protocols is not available in these BGP-only networks.

BGP Link-State (BGP-LS)[RFC9552] enables advertisement of a link state topology from link-state IGP protocols via BGP that can be consumed by a controller or in general any software component to get a complete topology view of the network. BGP-LS extensions for advertisement of certain aspects of a BGP topology for the Egress Peer Engineering (EPE) use-case [RFC9087] are specified in [RFC9086]. This document leverages the BGP-LS extensions that were defined for EPE and other BGP-LS features. The document specifies the procedures for advertising the underlying topology in a BGP-only network.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. BGP Routing in the Fabric

The applicability of this specification is limited to those deployments where BGP is used as hop-by-hop routing protocol between directly connected nodes in the fabric. While a data-center design [RFC7938] is used as a reference, the topology advertisement and its use for computation may also apply to other networks with BGP-only fabric or to BGP-only portions of a larger network topology.

BGP hop-by-hop routing can be setup using EBGp single-hop sessions over individual links between directly connected routers using their link addresses for peering as described in [RFC7938]. In such a design, the neighbors' link addresses may be provisioned for peering and the EBGp session operating directly over the link performs the monitoring of the neighbor on that link. A variation of this design would be that the EBGp session is setup between directly connected routers using their loopback sessions. The mechanisms for discovery of the neighbor's link addresses and their monitoring on a per link basis are outside the scope of this document.

Though this document uses the EBGp based design as a reference, it does not preclude other alternate designs using IBGP.

This document does not change base BGP routing protocol operations in the BGP-only network fabric that provides routing using the BGP best path selection process [RFC4271] .

3. Topology Collection Mechanism

To provide a topological view in networks where BGP is the only routing protocol, each BGP router advertises information about its local node, links, and prefixes. Figure 1 describes a typical deployment scenario. Every BGP router in the network is enabled for BGP-LS and forms BGP-LS sessions with one or more centralized BGP-LS speakers over which it sends its local topology information.

Each BGP router MAY also receive the topology information from all other BGP routers via these centralized BGP-LS speakers. This way, any BGP router (as also the centralized BGP-LS speakers) MAY obtain aggregated Link-State information for the entire BGP network. An external component (e.g. a controller) can obtain this information from the centralized BGP-LS speakers or directly by doing BGP-LS peering to the BGP routers. An internal software component on any of the BGP routers (e.g. TE module) can also receive the entire BGP network topology information from its local BGP process.




Figure 1: Link State Information Collection in a BGP Network

3.1. Peering Models

The peering model described above relies on the base BGP IPv4 or IPv6 routing underlay (e.g. as described in [RFC7938]) or any other mechanism for reachability for the BGP-LS session establishment with the centralized BGP speakers. A variation of this model would be to setup reachability to the centralized BGP speakers (or controller) over the out of band management network and for each BGP router in the fabric to use this management network for the BGP-LS session establishment with the centralized BGP speakers. This variation removes the dependency between the topology learning via BGP-LS from the reachability over the BGP routing in the fabric.

Another alternate design would be to enable the BGP-LS address-family as well on the hop-by-hop EBGp sessions in the underlay described in [RFC7938]. This approach results in the topology information being flooded via BGP-LS hop-by-hop along the BGP routers in the network. Other peering designs for BGP-LS sessions may also be possible and they are not precluded by this document.

4. Advertising BGP-only Network Topology

BGP-LS [RFC9552] defines the BGP-LS NLRI types (i.e. Node NLRI, Link NLRI and Prefix NLRI) along with their corresponding BGP-LS Attribute (i.e. Node Attribute, Link Attribute or Prefix Attribute) and the TLVs that map to the respective NLRI and Attribute for each type.

[RFC9086] specifies the BGP Protocol ID to be used for signaling BGP EPE information and the same is used for advertising of topology information in a BGP-only network.

[RFC9514] defines the BGP-LS NLRI that can be used to advertise Segment Routing for IPv6 (SRv6) Segment Identifier (SID) information instantiated on a BGP Router.

[RFC9857] defines the BGP-LS NLRIs that can be used to advertise information about Segment Routing (SR) Policies instantiated on a BGP Router headend.

The following sub-sections specify the use of these encodings by a router running BGP protocol.

4.1. Node Advertisements

[RFC9552] defines Node NLRI Type and the Node Descriptor TLVs as follows:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Protocol-ID |
+-----+-----+-----+-----+
|                               Identifier                               |
|                               (64 bits)                               |
+-----+-----+-----+-----+
//                               Local Node Descriptors (variable)       //
+-----+-----+-----+-----+

```

[RFC9086] introduces additional Node Descriptor TLVs for BGP protocol that are required to be used.

The following Node Descriptors TLVs MUST appear in the Node NLRI as Local Node Descriptors:

- * Autonomous System Number (TLV 512), which contains the advertising router ASN.
- * BGP Router-ID (TLV 516), which contains the BGP Identifier of the originating BGP router.

The BGP-LS Attribute associated with the Node NLRI MAY include the following TLVs that are defined in respective documents to signal the router properties and capabilities (Section 5.1 defines the procedures for their advertisements):

TLV Code Point	Description	Reference Document
266	Node MSD	[RFC8814]
1026	Node Name	[RFC9552]
1028	IPv4 TE Router-ID	[RFC9552]
1029	IPv6 TE Router-ID	[RFC9552]
1032	S-BFD Discriminators	[RFC9247]
1034	SR Capabilities	[RFC9085]
1035	SR Algorithm	[RFC9085]
1036	SR Local Block	[RFC9085]
1038	SRv6 Capabilities	[RFC9514]

Table 1: Node Attribute TLVs

The above list of TLVs is not exhaustive but indicative as of the time of writing of this document.

4.2. Link Advertisements

[RFC9552] defines Link NLRI Type and its Node and Link Descriptor TLVs as follows:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Protocol-ID |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Identifier                               |
|                               (64 bits)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                               Local Node Descriptors (variable)       //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                               Remote Node Descriptors (variable)      //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                               Link Descriptors (variable)             //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The following Node Descriptors TLVs MUST appear in the Link NLRI as Local Node Descriptors:

- * Autonomous System Number (TLV 512), which contains the advertising router ASN.
- * BGP Router-ID (TLV 516), which contains the BGP Identifier of the originating BGP router.

The following Node Descriptors TLVs MUST appear in the Link NLRI as Remote Node Descriptors:

- * Autonomous System Number (TLV 512), which contains the peer ASN.
- * BGP Router-ID (TLV 516), which contains the BGP Identifier of the peer BGP router.

The following Link Descriptors TLVs MUST appear in the Link NLRI as Link Descriptors:

- * Link Local/Remote Identifiers (TLV 258) containing the 4-octet Link Local Identifier followed by the 4-octet Link Remote Identifier. The value 0 MUST be used for the Link Remote Identifier when the value is unknown.

In addition, the following Link Descriptors TLVs SHOULD appear in the Link NLRI as Link Descriptors based on the address family used for setting up the BGP Peering or the addresses configured on the links:

- * IPv4 Interface Address (TLV 259) contains the address of the local interface through which the BGP session is established using IPv4 address.
- * IPv4 Neighbor Address (TLV 260) contains the IPv4 address of the peer interface used by the BGP session establishment using IPv4 address.
- * IPv6 Interface Address (TLV 261) contains the address of the local interface through which the BGP session is established using IPv6 address.
- * IPv6 Neighbor Address (TLV 262) contains the IPv6 address of the peer interface used by the BGP session establishment using IPv6 address.

The BGP-LS Attribute associated with the Link NLRI MAY include the following TLVs that are defined in respective documents to signal the router's local links' properties and capabilities (Section 5.2 defines the procedures for their advertisements) :

TLV Code Point	Description	Reference Document
267	Link MSD	[RFC8814]
1088	Administrative group (color)	[RFC9552]
1089	Maximum link bandwidth	[RFC9552]
1092	TE Default Metric	[RFC9552]
1096	SRLG	[RFC9552]
1098	Link Name	[RFC9552]
1101	EPE Peer Node SID	[RFC9086]
1102	EPE Peer Adj SID	[RFC9086]
1103	EPE Peer Set SID	[RFC9086]
1106	SRv6 End.X SID	[RFC9514]
1114	Unidirectional link delay	[RFC8571]
1115	Min/Max Unidirectional link delay	[RFC8571]
1116	Unidirectional delay variation	[RFC8571]
1117	Unidirectional link loss	[RFC8571]
1172	L2 Bundle Member	[RFC9085]
1173	Extended Administrative group (color)	[RFC9104]

Table 2: Link Attribute TLVs

The above list of TLVs is not exhaustive but indicative as of the time of writing of this document.

4.3. Prefix Advertisements

[RFC9552] defines Prefix NLRI Type and its Node and Prefix Descriptor TLVs as follows:

```

      0                      1                      2                      3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Protocol-ID |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Identifier
|                                     (64 bits)
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Local Node Descriptors (variable)                               //
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Prefix Descriptors (variable)                               //
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The following Node Descriptors TLVs MUST appear in the Prefix NLRI as Local Node Descriptors:

- * Autonomous System Number (TLV 512), which contains the advertising router ASN.
- * BGP Router-ID (TLV 516), which contains the BGP Identifier of the originating BGP router

The Prefix Descriptor MUST contain the IP Reachability Information TLV (TLV 265) to identify the prefix.

This document defines a new BGP Route Type TLV that MUST be included in the Prefix Descriptor when the BGP node advertises the Prefix NLRI. The format of this TLV is as follows:

```

      0                      1                      2                      3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Type                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Route Type |                                     Length                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Where:

Type: 2-octet field with value TBD.

Length: 2-octet field with value set to 1.

Route Type: 1-octet with the following values defined:

Value	Type	Description
1	Local	Local interface prefix e.g. Loopback
2	Attached	Directly attached node's prefix e.g host
3	External BGP	Prefix learnt via EBGp
4	Internal BGP	Prefix learnt via IBGP
5	Redistributed	Prefix redistributed into BGP

Figure 2: BGP Route Types

The BGP-LS Attribute associated with the Prefix NLRI MAY include the following TLVs that are defined in respective documents to signal the router's own prefix properties and capabilities (Section 5.3 defines the procedures for their advertisements):

TLV Code Point	Description	Reference Document
1155	Prefix Metric	[RFC9552]
1158	Prefix SID	[RFC9085]
1162	SRv6 Locator	[RFC9514]
1170	Prefix Attributes Flags	[RFC9085]
1171	Source Router Identifier	[RFC9085]

Table 3: Prefix Attribute TLVs

The above list of TLVs is not exhaustive but indicative as of the time of writing of this document.

4.4. SR Policy Advertisements

[RFC9857] defines SR Policy Candidate Path NLRI Type with its Headend Node and SR Policy Candidate Path Descriptor TLVs as follows:

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Protocol-ID |
+-----+-----+-----+-----+
|                               Identifier                               |
|                               (64 bits)                               |
+-----+-----+-----+-----+
//                               Headend (Node Descriptors)                               //
+-----+-----+-----+-----+
//          SR Policy Candidate Path Descriptors (variable)          //
+-----+-----+-----+-----+

```

The Headend Node Descriptors TLVs are the same as specified in Section 4.1. The semantics for the SR Policy Candidate Path Descriptor TLVs and the TLVs associated with the BGP-LS Attribute are used as specified in [RFC9857].

4.5. SRv6 SID Advertisements

[RFC9514] defines SRv6 NLRI Type and its Local Node and SRv6 SID Descriptor TLVs as follows:

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
| Protocol-ID |
+-----+-----+-----+-----+
|                               Identifier                               |
|                               (8 octets)                               |
+-----+-----+-----+-----+
|                               Local Node Descriptors (variable)                               //
+-----+-----+-----+-----+
|                               SRv6 SID Descriptors (variable)                               //
+-----+-----+-----+-----+

```

The Local Node Descriptors TLVs are the same as specified in Section 4.1. The semantics for the SRv6 SID Descriptor TLVs and the TLVs associated with the BGP-LS Attribute are used as specified in [RFC9514].

5. Procedures

In a network where BGP is the only routing protocol, the BGP-LS session is used to advertise the necessary information about the local node properties, its local links' properties and where necessary the prefix's owned by the node. TE Policies, that are instantiated on the local node (i.e. when it is the head-end for the policy), along with their properties are also advertised via the BGP-LS session. This information, once collected across all BGP routers in the network, provides a complete topology view of the network. Many of these attributes are not part of the base BGP protocol operations and are either configured or provided by other components on the router. This information needs to be collected from within the node and advertised out via BGP-LS.

The following sections describe the procedures for the propagation of the BGP-LS NLRIs on a BGP router into the BGP-LS session. These procedures for propagation of BGP topology information via BGP-LS SHOULD be applied only in deployments and use-cases where necessary and SHOULD NOT be applied in every BGP deployment when BGP-LS is enabled. Implementations MAY provide a configuration option to enable these procedures in required deployments.

5.1. Advertisement of Router's Node Attributes

Advertisement of the Node NLRI via BGP-LS by each BGP router in a BGP-only network enables the discovery of all the router nodes in the topology. The Node NLRI MUST be generated by a BGP router only for itself and even when there are no attributes to be advertised along with it.

The Node attributes defined currently related to Segment Routing (SR) [RFC8402] have been described in Table 1 and are to be advertised when SR is enabled. This includes:

- * All SR enabled routers support the default SR algorithm 0 and MUST advertise it in the SR Algorithm TLV. Other algorithms (including Flexible Algorithm [RFC9350]) SHOULD be advertised when supported.
- * The Segment Routing Global Block (SRGB) provisioned on the router which is used by BGP Prefix SIDs [RFC8669] and other SR control plane protocols on the router MUST be advertised. The value for Flags field in the TLV is not defined for BGP protocol and MUST be set to 0 by the originator and ignored by receivers.

- * The Segment Routing Local Block (SRLB) provisioned on the router which MAY be used by BGP EPE SIDs [RFC9086] SHOULD be advertised. The value for Flags field in the TLV is not defined for BGP protocol and MUST be set to 0 by the originator and ignored by receivers.
- * The Node level MSD provides the Node's capabilities for SR SID operations and SHOULD be advertised.
- * When the router supports SR Flexible Algorithms and is provisioned with the Flexible Algorithm Definition (FAD), then it MUST advertise the same.

The Node Name Attribute SHOULD be advertised when available.

This document introduces some of the TE concepts into BGP-only networks. Provisioning of TE Router-ID with a unique address normally associated with a loopback interface on the router enables TE use-cases for both IPv4 and IPv6 SHOULD be supported. The BGP Router-ID along with the ASN also provides the capability for uniquely identifying a BGP router in the network.

Other Node Attributes applicable to a BGP Router may also be included and this document does not describe the exhaustive list.

5.2. Advertisement of Router's Local Links Attributes

Each BGP router in a BGP-only network also advertises its local links using the Link NLRIs thru BGP-LS. The Link NLRI for a given link between two BGP routers is advertised as uni-directional logical "half-link" and its link descriptors allow the correlation between the two NLRIs "half-links" originated by the peering routers to describe the bi-directional logical link and its attributes on both routers.

The discovery of all the links and their local and remote identifiers in a BGP-only network relies on the design that uses EBGp sessions over each interconnecting link using the link IP addresses (refer [RFC7938]). In this case, a Link NLRI MUST be generated by a BGP router for each of its local link regardless of whether it has any link attributes to be advertised for it.

When doing EBGp multi-hop sessions between directly connected BGP routers, the underlying link information would need to learn by some discovery protocol or provisioning entity. The mechanisms to learn the underlying link information for BGP-LS advertisements are outside the scope of this document. However, to provide a true link topology picture, the advertisement of underlying links is RECOMMENDED for most use-cases instead of a single EBGp peering representation of a link between the routers using their loopback addresses.

The Link NLRI represents an adjacency between BGP routers and its association with the underlying Layer 3 link. When the underlying Layer 3 link or the BGP session on top of it goes down, the Link NLRI MUST be withdrawn by the BGP router. The monitoring of links, detecting of their failures and notification to BGP may be performed using mechanisms like BFD. This enables faster detection of failures and verification of the underlying links.

Advertisement of the Link NLRIs via BGP-LS by each BGP router in a BGP-only network enables the discovery of all the active links in the topology.

TE attributes for links have been traditionally associated with Link State Routing protocols. However, with the ability to discover the link topology via BGP-LS as specified in this document, the TE attributes and their principles can also be applied to a network running BGP alone. The TE attributes for a link have been described in Table 2 and MAY be advertised when TE use-cases are enabled. This includes:

- * The maximum bandwidth of a link is its protocol independent attribute and SHOULD be advertised.
- * TE concepts of Administrative Groups (also known as affinities) and Shared Risk Link Groups (SRLGs) MAY be provisioned locally on links and then MUST be advertised.
- * The BGP base protocol does not operate with link metrics, however, a TE metric concept can be introduced in a BGP only network as well for TE use-cases. Implementations MAY provide the ability to provision TE metric value for a link for BGP use including a different default value for it. The TE metric attribute SHOULD be advertised for each link when configured and its default value is taken as 100. When not advertised for a link, implementations who intend to use the TE metric MUST assume the value to be 100.

- * The delay and loss TE metrics for links are measured via MPLS Performance Monitoring [RFC6374] and their measurement mechanism over a link are independent of the routing protocol. The same mechanism MAY be enabled in BGP-only networks and their values advertised via BGP-LS.

The Link attributes defined currently related to the Segment Routing feature BGP EPE [RFC9086] have been described in Table 2 and are to be advertised when SR use-cases are enabled. This includes:

- * The BGP Peering SIDs provide a functionality similar to Adjacency-SID (refer [RFC8402]) in BGP-only networks. Implementations SHOULD allocate the BGP Peer-Adjacency SID for all its links and the BGP Peer-Node SID for all its peer routers. Implementations MAY allocate the BGP Peer-Set SID based on local configuration.
- * The Link level MSD provides the per link capabilities for SR SID operations and SHOULD be advertised when the router links have differing capabilities.

The use of Layer 3 bundle links which comprise of multiple layer 2 member links are often used in BGP networks. When BGP session is configured over such a layer 3 link, the link attributes of the underlying layer 2 links MAY be advertised individually using the L2 Bundle Member TLV. The applicable attributes for the L2 links are described in [RFC9085].

The Link Name Attribute MAY be advertised when available.

Other Link Attributes applicable to a BGP Router may also be included and this document does not describe the exhaustive list.

5.3. Advertisement of Router's Prefix Attributes

Advertisement of the Prefix NLRI via BGP-LS may be required only in specific use-cases. Since the base BGP protocol along with its extensions already signals Prefix reachability via different NLRIs, there is no necessity to duplicate the information via BGP-LS session. However, for specific use-cases related to SR Traffic Engineering (SR-TE), it is required for each router to advertise its Prefix SID(s) (refer [RFC8402]) that can be used to direct traffic via specific BGP routers. Advertising such BGP Prefix SID for every BGP router provides this key attribute via BGP-LS and avoids the requirement for the consumer of the topology information (e.g. a controller or local TE process) to tap into other BGP NLRI information.

Advertisement of the Prefix NLRI via BGP-LS MUST be done for its locally configured prefixes (e.g. its loopback interface address) and when BGP is advertising the BGP Prefix SID ([RFC8669]) for it. The advertisement of the Prefix NLRI via BGP-LS for other prefixes learnt by the router MAY be done based on the specific use-case requirement and the BGP Route Type as described in Figure 2 indicates the type of route being advertised.

The Prefix attributes defined currently related to SR [RFC8402] have been described in Table 3 and MAY be advertised when SR is enabled. This includes:

- * The Prefix SID TLV is included with the SID advertised as the index to be consistent with the Label-Index TLV of BGP Prefix SID attribute. The default algorithm is MUST be set to 0 by the originator except in the case where a local prefix is associated with a specific SR Algorithm. The flags are defined as the most significant 8 bits of the 16 bit field defined for Label-Index TLV in [RFC8669].
- * For certain SR-TE uses, the Prefix Metric value MAY be included and it is set based on the SR-TE computation based on the link-state topology learnt via BGP-LS.

Other Prefix Attributes applicable may also be included and this document does not describe the exhaustive list.

5.4. Advertisement of Router's SR Policy Candidate Path Attributes

SR Policies that are setup using SR-TE mechanisms MAY be instantiated on a BGP router. One use-case that results in such SR Policy instantiation on a BGP router is described later in this document in Section 6.2. Advertising such SR Policy Candidate Paths instantiated for every BGP router as head-end via BGP-LS provides the consumer of the topology information (e.g. a controller or local TE process) a policy view of the BGP fabric as well.

The procedures for advertisement of the SR Policy Candidate Path NLRI via BGP-LS MUST be done only for its locally instantiated SR Policies and as specified in [RFC9857]. Implementation MAY provide configuration options to control the specific set of SR Policies that are to be advertised from the local node.

5.5. Advertisement of Router's SRv6 SID Attributes

The SRv6 End SID instantiated on a BGP router can be used to signal SRv6 capabilities for the supported services. The advertisement of the SRv6 SID NLRI via BGP-LS may be required on SRv6 capable routers.

The SRv6 SID attributes have been described in [RFC9514] and MAY be advertised when SRv6 is enabled. This includes:

- * The SRv6 Endpoint Behavior defines specific behaviors for the SRv6 SID and must be advertised.

6. Usage of BGP Topology

This section describes some of the use-cases for the building of the BGP topology information as specified in this document and leveraging it for enabling new functionality.

6.1. Topology View for Monitoring

The BGP-LS advertisement of the BGP topology as specified in this document provides a live topology view of the BGP network for an application or controller that is monitoring the network. The topology view is from the BGP protocol perspective and includes the underlying links as well that aids in network monitoring as well as diagnostics use-cases. BGP-LS is the de-facto protocol for northbound propagation of network topology related information for most IGP networks and extending this capability for BGP-only networks allows existing controllers and applications to consume the information with some incremental BGP protocol awareness.

6.2. SR-TE in BGP Networks

The SR-TE use-case for BGP builds on top of functionality specified in [RFC8669] and also described in [RFC8670]. The BGP SR Prefix SID signaled, provides the basic connectivity between all BGP routers using their loopback addresses. This provides the basic best-effort paths in the network using the base BGP decision process that is unchanged. BGP and other overlay routes and services recurse on top of these loopback addresses of the egress nodes and the forwarding is done via the BGP SR Prefix SID labels in the underlay. While this version of the document focuses on the examples with MPLS dataplane instantiation for SR, the same is applicable for the IPv6 dataplane instantiation (SRv6) as well.

SR-TE for BGP provides underlay paths through the network for the overlay routes and services with specific SLA requirements and use-cases like path disjointness, low latency paths, inclusion or exclusion and other TE considerations.

[RFC9256] specifies the SR Policy architecture. [RFC9830] and [RFC9831] describes the extensions to BGP for signaling of SR Policies from a controller to the SR-TE headend BGP router. BGP-LS has been extended to allow signaling of the SR Policies from SR-TE

head-end to controllers via [RFC9857] which allows the controllers to learn the state of SR Policies instantiated on routers in the network. This document completes the missing piece that is related to getting the BGP topology information from all the routers to a controller as well the local SRTE process on each router for their path computation requirements.

The signaling of SR Policies from controller to SR-TE headend and reporting of the state back to the controller can also be done using PCEP ([RFC8664], [RFC8281], [RFC8231]). However, the BGP topology learning via BGP-LS which is specified in this document is also required for the deployments that uses PCEP in the BGP-only network.

The topology collected via BGP-LS in a BGP-only fabric in a Segment Routing deployment comprise of:

- * The properties of every BGP router node and the Prefix SIDs to reach that node.
- * The properties of all the links between the BGP routers and the Peer-Adjacency-SIDs (and other EPE SIDs) corresponding to them that allow directing traffic over specific links and/or to specific neighbors.
- * The properties and state of the SR Policies instantiated on each of the BGP routers along with their end points, their properties and most importantly the Binding SID to steer traffic into the SR Policies.

This topology information allows a computation node to build SR Policies for services over the BGP fabric for a given traffic engineering objective at any given node.

The topology of the BGP fabric is advertised to a centralized controller or application for use-cases that need a centralized computation of SR Policy which can then be signaled to the SR-TE head-end node via PCEP or BGP-SRTE. The topology may also be distributed to any node in the BGP fabric to be used by its local SR-TE process to perform path computation for its own SR Policies for use-cases that are addressed by local computation.

A high level summary of the key topology information advertised via BGP-LS by BGP routers can be used for TE computations as follows

- * The BGP SR Prefix SIDs and the BGP EPE Peering Adjacency SIDs provide the equivalent of the IGP Prefix and Adjacency SIDs and can be used to direct traffic to a specific BGP router and over a specific BGP peer session or link respectively. Traffic for the BGP SR Prefix SIDs follow the path computed by the BGP decision process.
- * The TE metric can be used to tailor the choice of specific paths in the network for SR-TE.
- * The TE administrative group (also known as affinities) and SRLG attributes can be configured over links to enable computation of paths with inclusion and exclusion of specific links or paths that are mutually disjoint.
- * The enabling of link delay and loss measurements and their advertisements can help monitoring the link quality and carve out paths based on latency and other SLA requirements.
- * The signaling of the Node and Link MSD allows controllers to instantiate SR Policies based on the capability of the routers.

This section attempts to highlight and describe at a high level some of the possible SR-TE solutions and use-cases in a BGP-only network. The actual SR-TE computation and algorithms are outside the scope of this document.

7. IANA Considerations

This document requests IANA to allocate code points from the "BGP-LS NLRI and Attribute TLVs" registry of the "Border Gateway Protocol - Link-State (BGP-LS) Parameters" registry group.

This document requests the allocation following TLV codepoints:

TLV Code Point	Description	Reference
TBD	BGP Route Type	this document

8. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [RFC4271] for a discussion of BGP security. Also refer to [RFC4272] and [RFC6952] for analysis of security issues for BGP.

[RFC9552] defines BGP-LS NLRI to which the extensions defined in this document apply. Section 10 of [RFC9552] also applies to these extensions. The procedures defined in this document, by themselves, do not affect the BGP-LS security model discussed in [RFC9552].

The BGP-LS extensions specified in this document enable topology visibility and traffic engineering use-cases within a BGP-only fabric as described in this document. BGP-LS operates within a trusted domain and its security considerations apply to BGP sessions when carrying topology information. The topology information distributed by BGP-LS is expected to be used entirely within this trusted domain which comprises a single AS or multiple ASes/domains within a single provider network. Therefore, precaution is necessary to ensure that the topology information advertised via BGP-LS sessions is limited to nodes in a secure manner within this trusted domain.

Additionally, it should be considered that the export of detailed topology information, as described in this document, constitutes a risk to confidentiality of mission-critical or commercially sensitive information about the network. BGP-LS peerings are not automatic and require configuration; thus, it is the responsibility of the network operator to ensure that only trusted nodes (that include both routers and controller applications) within the trusted domain are configured to receive such information.

9. Acknowledgements

The authors would like to thank Bruno Decraene for his review and comments on this document.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.
- [RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.
- [RFC8814] Tantsura, J., Chunduri, U., Talaulikar, K., Mirsky, G., and N. Triantafyllis, "Signaling Maximum SID Depth (MSD) Using the Border Gateway Protocol - Link State", RFC 8814, DOI 10.17487/RFC8814, August 2020, <<https://www.rfc-editor.org/info/rfc8814>>.
- [RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <<https://www.rfc-editor.org/info/rfc9085>>.
- [RFC9086] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Patel, K., Ray, S., and J. Dong, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering", RFC 9086, DOI 10.17487/RFC9086, August 2021, <<https://www.rfc-editor.org/info/rfc9086>>.
- [RFC9104] Tantsura, J., Wang, Z., Wu, Q., and K. Talaulikar, "Distribution of Traffic Engineering Extended Administrative Groups Using the Border Gateway Protocol - Link State (BGP-LS)", RFC 9104, DOI 10.17487/RFC9104, August 2021, <<https://www.rfc-editor.org/info/rfc9104>>.

- [RFC9247] Li, Z., Zhuang, S., Talaulikar, K., Ed., Aldrin, S., Tantsura, J., and G. Mirsky, "BGP - Link State (BGP-LS) Extensions for Seamless Bidirectional Forwarding Detection (S-BFD)", RFC 9247, DOI 10.17487/RFC9247, June 2022, <<https://www.rfc-editor.org/info/rfc9247>>.
- [RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.
- [RFC9514] Dawra, G., Filsfils, C., Talaulikar, K., Ed., Chen, M., Bernier, D., and B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)", RFC 9514, DOI 10.17487/RFC9514, December 2023, <<https://www.rfc-editor.org/info/rfc9514>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.
- [RFC9857] Previdi, S., Talaulikar, K., Ed., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies Using BGP - Link State", RFC 9857, DOI 10.17487/RFC9857, October 2025, <<https://www.rfc-editor.org/info/rfc9857>>.

10.2. Informative References

- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <<https://www.rfc-editor.org/info/rfc6374>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <<https://www.rfc-editor.org/info/rfc6952>>.

- [RFC7938] Lapukhov, P., Premji, A., and J. Mitchell, Ed., "Use of BGP for Routing in Large-Scale Data Centers", RFC 7938, DOI 10.17487/RFC7938, August 2016, <<https://www.rfc-editor.org/info/rfc7938>>.
- [RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/info/rfc8281>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [RFC8670] Filsfils, C., Ed., Previdi, S., Dawra, G., Aries, E., and P. Lapukhov, "BGP Prefix Segment in Large-Scale Data Centers", RFC 8670, DOI 10.17487/RFC8670, December 2019, <<https://www.rfc-editor.org/info/rfc8670>>.
- [RFC9087] Filsfils, C., Ed., Previdi, S., Dawra, G., Ed., Aries, E., and D. Afanasiev, "Segment Routing Centralized BGP Egress Peer Engineering", RFC 9087, DOI 10.17487/RFC9087, August 2021, <<https://www.rfc-editor.org/info/rfc9087>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.
- [RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Ed., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025, <<https://www.rfc-editor.org/info/rfc9830>>.

[RFC9831] Talaulikar, K., Ed., Filsfils, C., Previdi, S., Mattes, P., and D. Jain, "Segment Type Extensions for BGP Segment Routing (SR) Policy", RFC 9831, DOI 10.17487/RFC9831, September 2025, <<https://www.rfc-editor.org/info/rfc9831>>.

Authors' Addresses

Ketan Talaulikar
Cisco Systems
India
Email: ketant.ietf@gmail.com

Aravind Babu MahendraBabu (editor)
Cisco Systems
India
Email: aramahen@cisco.com

Clarence Filsfils
Cisco Systems
Brussels
Belgium
Email: cfilsfil@cisco.com

Krishna Swamy
Cisco Systems
San Jose,
United States of America
Email: kriswamy@cisco.com

Shawn Zandi
LinkedIn
United States of America
Email: szandi@linkedin.com

Gaurav Dawra
LinkedIn
United States of America
Email: gdawra.ietf@gmail.com

Muhammad Durrani
Equinix
United States of America

Email: mdurrani@equinix.com